Modular Verification of the Evolving Internet

Pamela Zave and Jennifer Rexford

Princeton University

Great Momentum for Network Verification

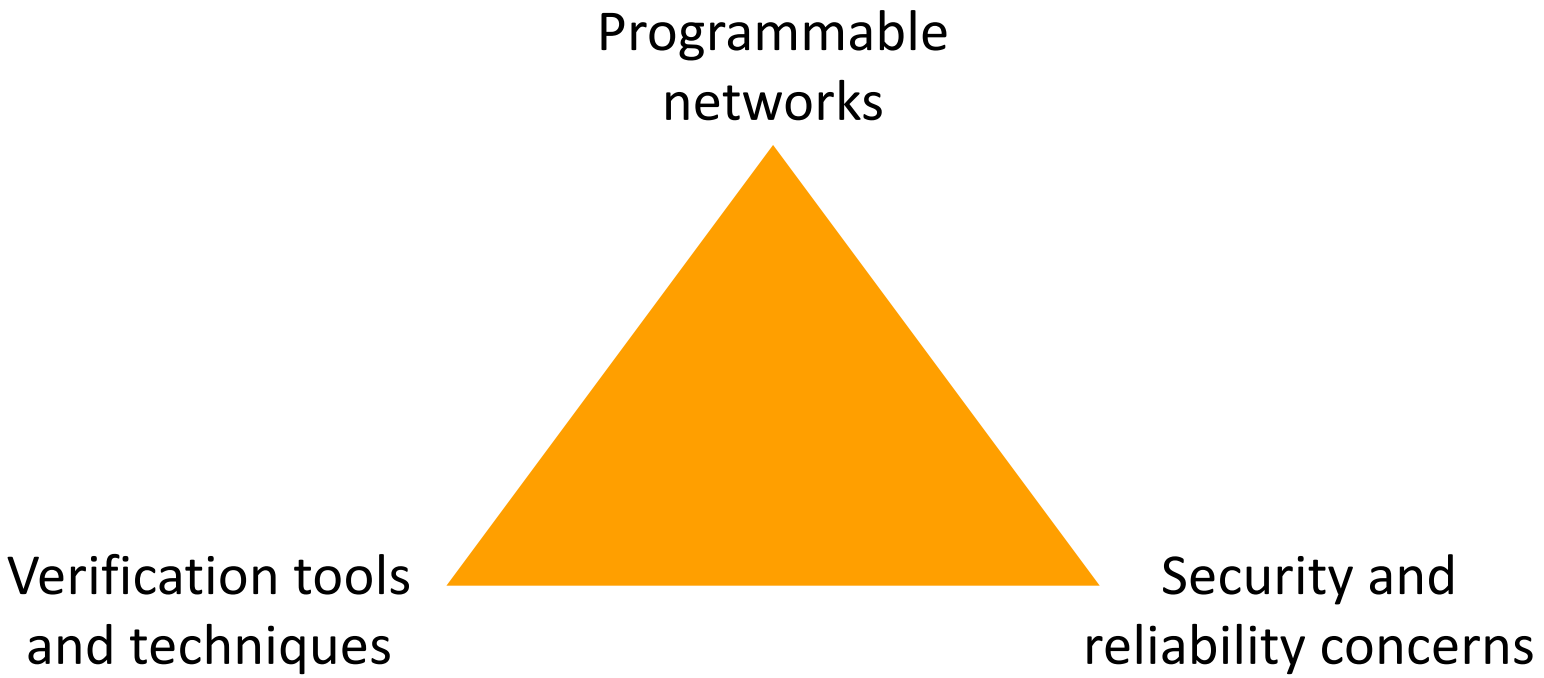

Success Verifying Low-Level Properties

- Verifying low-level properties
  - Of protocols (e.g., TCP, QUIC, BGP)
  - Of router configurations (e.g., BGP policies)
  - Of data-plane state (e.g., match-action tables)
  - Of data-plane programs (e.g., P4 code)
- E.g., basic reachability properties
  - No forwarding loops
  - No (unintended) blackholes
  - Filtering of unwanted packets
- Using a variety of verification techniques

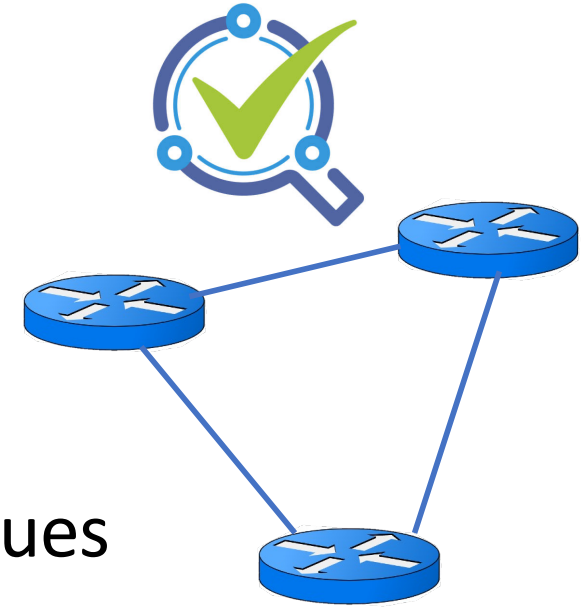

Verify Communication *Services*

- Network provides communication services
  - For the components of a distributed system
  - E.g., session affinity, middlebox insertion, confidentiality
  - … so, we should verify these important properties
- Data-plane state is *not* the right abstraction
  - Incomplete, indirect reflection of services
  - Misses the modular structure of the services

- Verification limitations
  - Detecting problems too late (reactive)
  - Hard to determine the right place(s) to fix
  - Poor scalability of the verification process

Verification With Broader Scope

- Interaction with other system components
  - The network rarely acts alone
  - E.g., compute, storage, databases, etc.

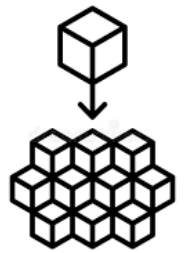

- Interaction with the rest of the Internet
  - Protocols: routing and transport protocols
  - Security: VPNs, certificate authorities
  - Middleboxes: NAT, firewall, load balancer

Toward Verifying Networked Systems

Understand big picture

See where system fits in the larger whole

… and properties it needs to fit properly

Strive for modularity

For ease of understanding

To enable scalable verification

To leverage right tools in right places

Ensure pieces fit together

Adopt a compositional model

An employee is connecting to the personnel database of his/her enterprise . . .

. . . and we want to prove that this communication is secure.

Yes! There is a packet filter/virus scanner in the enterprise . . .

. . . but the enterprise's network has to route to middleboxes with session affinity . . .

. . . and the employee is not in the office, but sitting in a coffee shop . . .

. . . and the employee is using a VPN . . .

. . . and the coffee shop's network has private IP addresses and a NAT . . .

How can we understand how all these pieces fit together?

THE CLASSIC INTERNET ARCHITECTURE

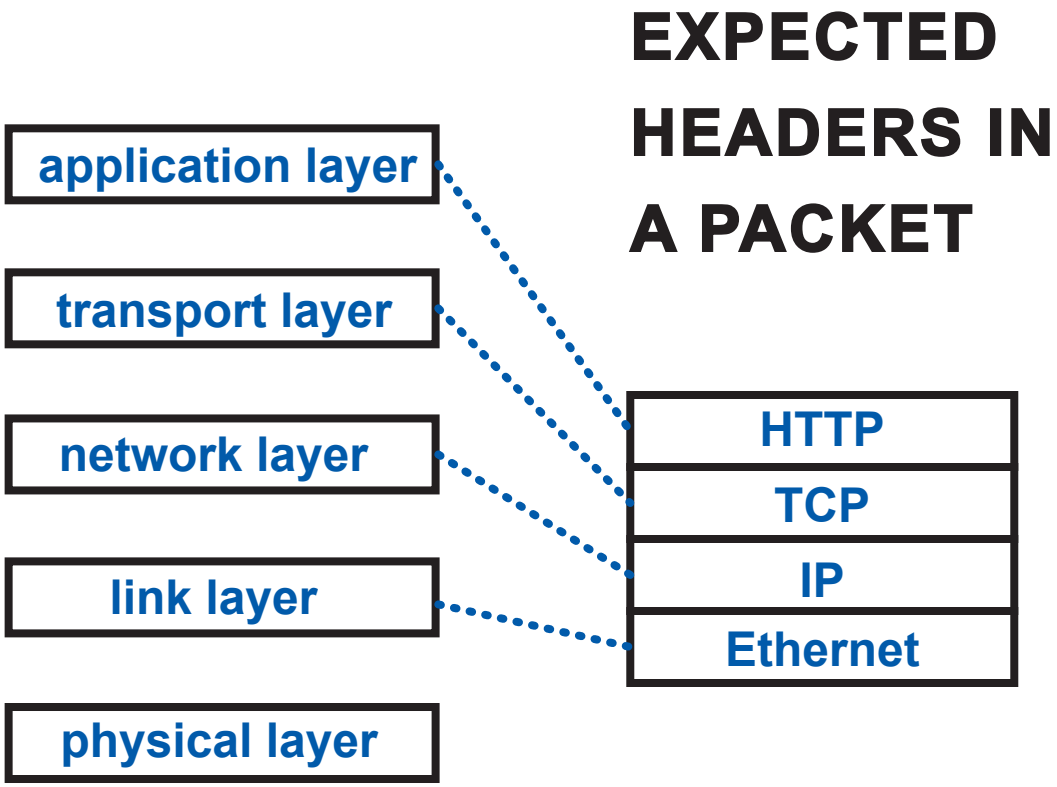

the Internet architecture is the IP protocol suite . . . and it has not changed significantly since 1993

so the Internet architecture has not evolved, although it should!
(but we don't know how)

THE CLASSIC INTERNET ARCHITECTURE

AN EXAMPLE OF THE REAL INTERNET ARCHITECTURE

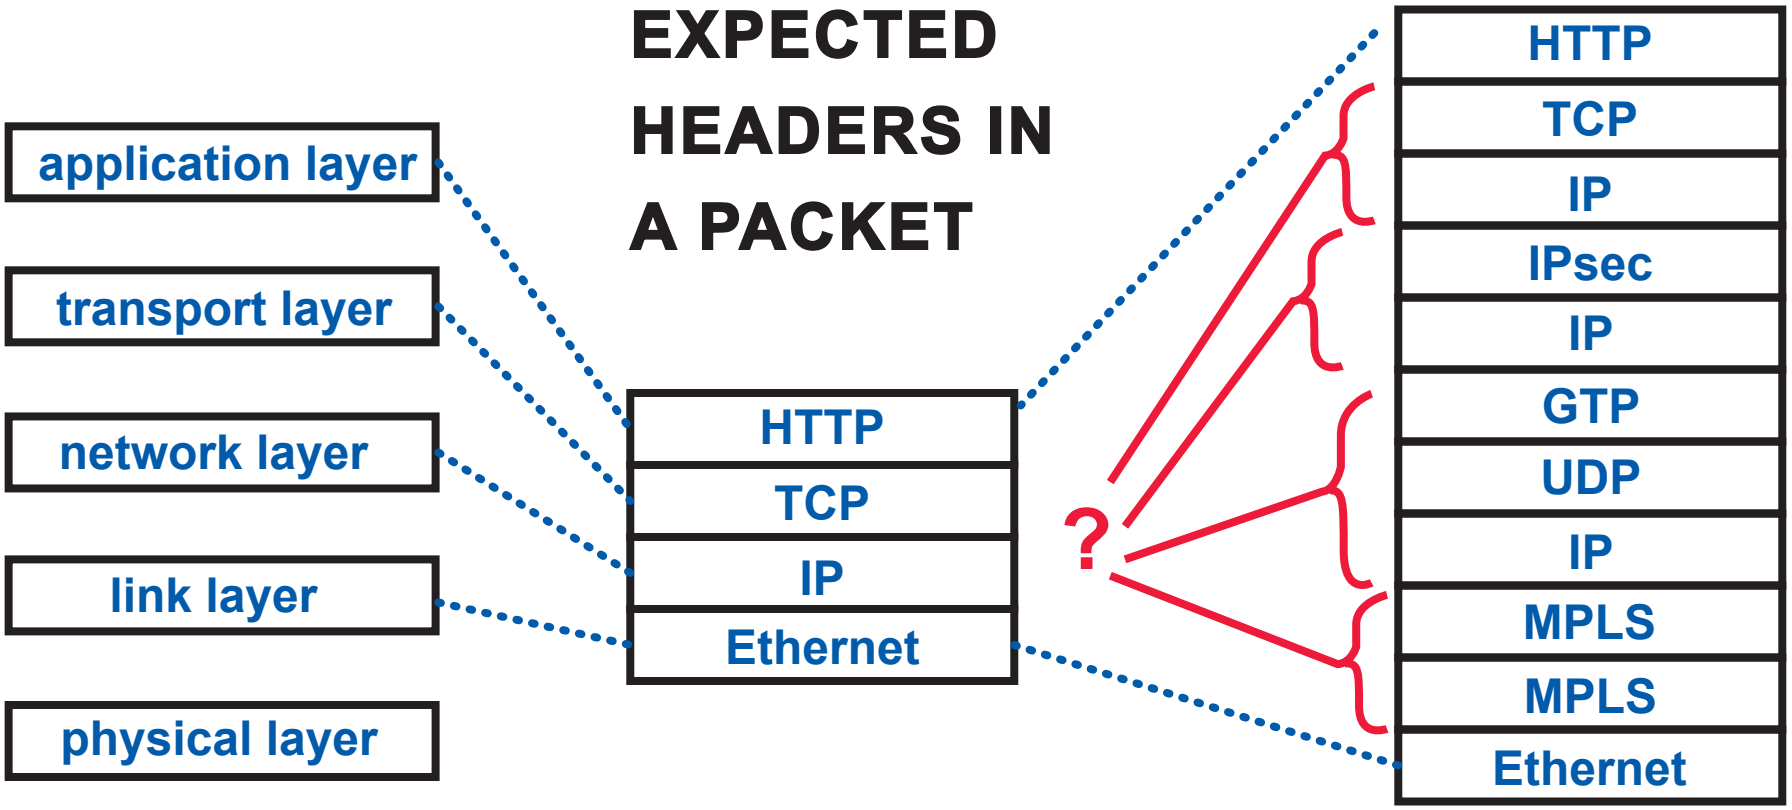

the Internet architecture is the IP protocol suite . . . and it has not changed significantly since 1993

so the Internet architecture has not evolved, although it should!
(but we don't know how)

the Internet architecture has evolved—to this . . .

. . . but the classic architecture gives us no way to talk about it or understand it

COMPOSITIONAL NETWORK ARCHITECTURE

A NEW MODEL OF NETWORK ARCHITECTURE

GENERAL: describes the architecture of the Internet—

past, present, and many possible futures

ACCURATE AND PRECISE

- carefully-chosen terminology
- much of the model has been formalized

REVEALS THE INHERENT MODULARITY IN NETWORK ARCHITECTURES

- separation of concerns
- recognition of recurring patterns
- modular verification

the old End-to-End Principle
is dead,
but now we have a new one!

THE MODULE IS A NETWORK:

IN OUR TERMINOLOGY, A SELF-CONTAINED MICROCOSM OF NETWORKING

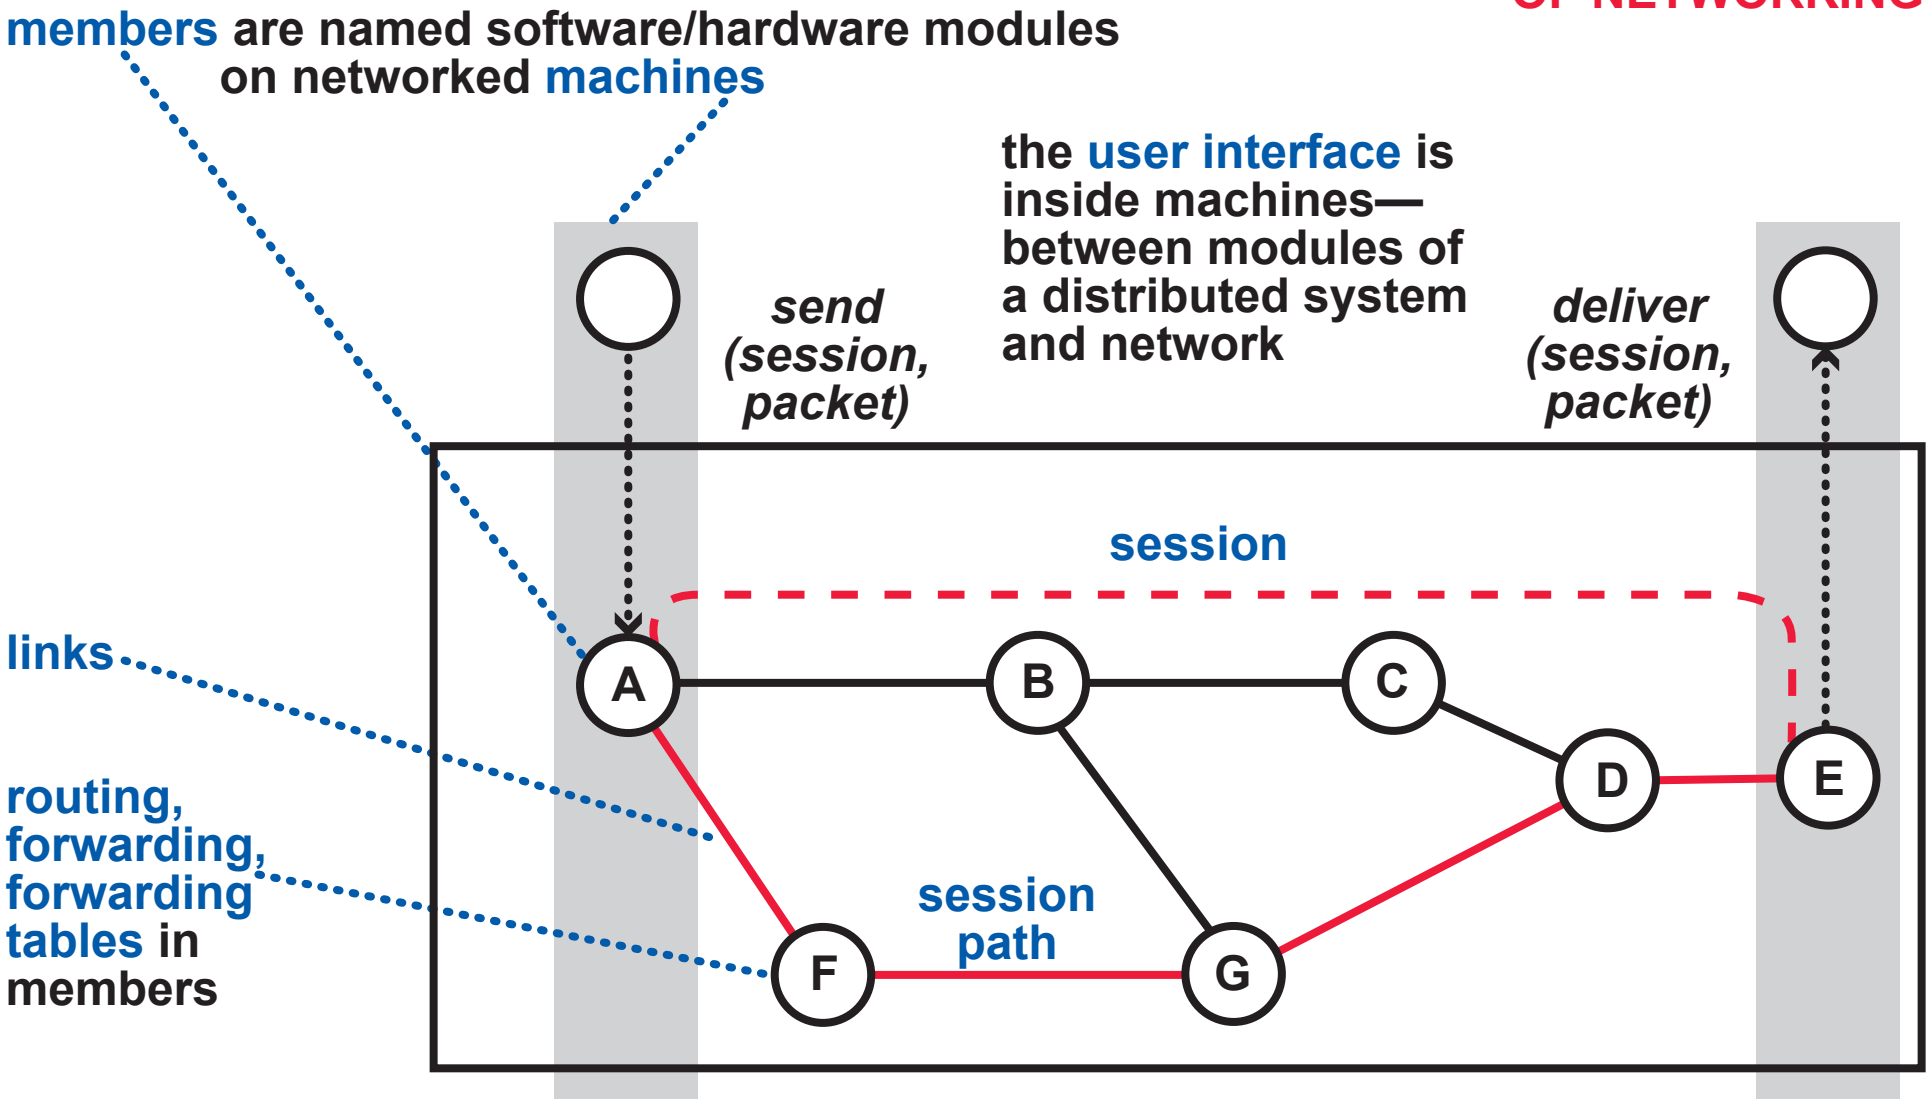

session protocols are part of the network design

sessions are usages of the network—central to its state and behavior

COMPOSITION OPERATORS ON NETWORKS 1

BRIDGING: The autonomous IP networks of the Internet are composed in this way.

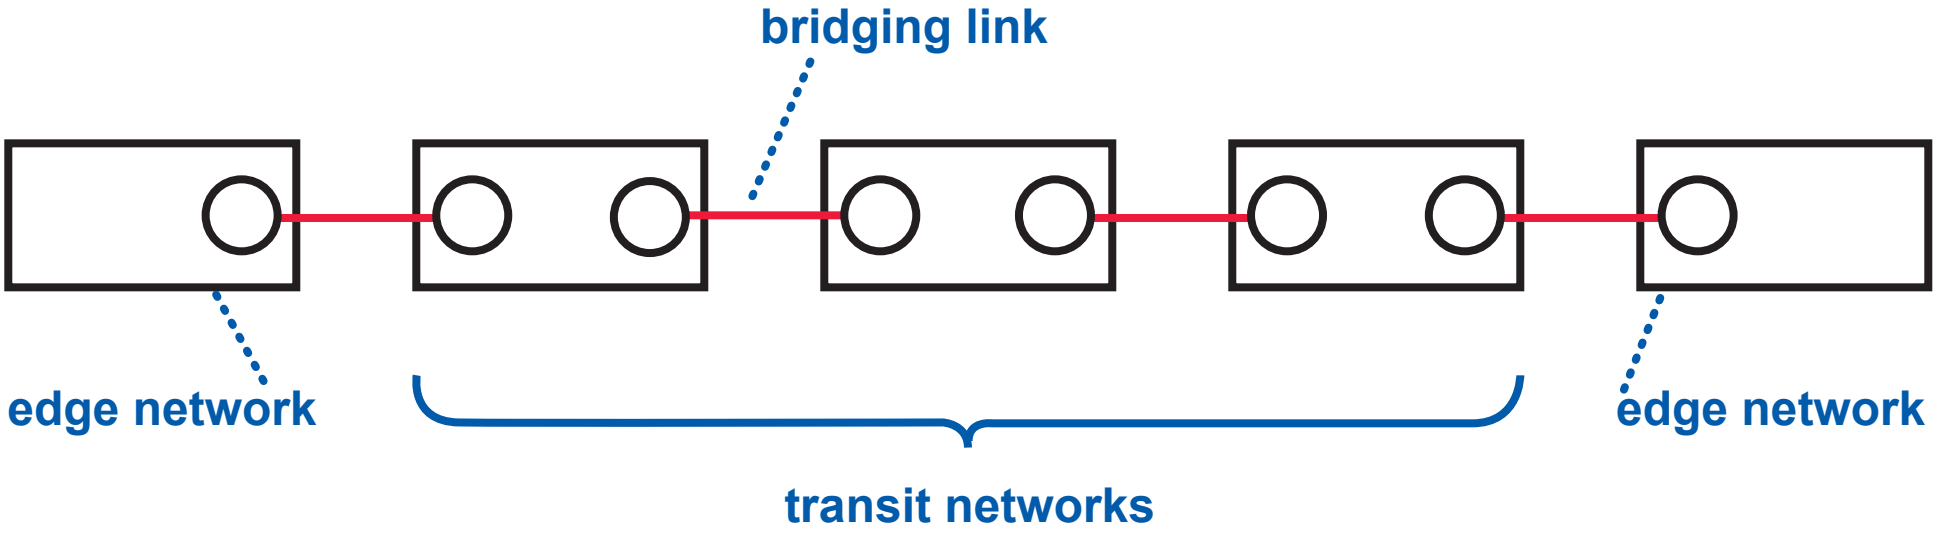

COMPOSITION OPERATORS ON NETWORKS 2

LAYERING:
A link of an overlay network is implemented by a session of an underlay network.

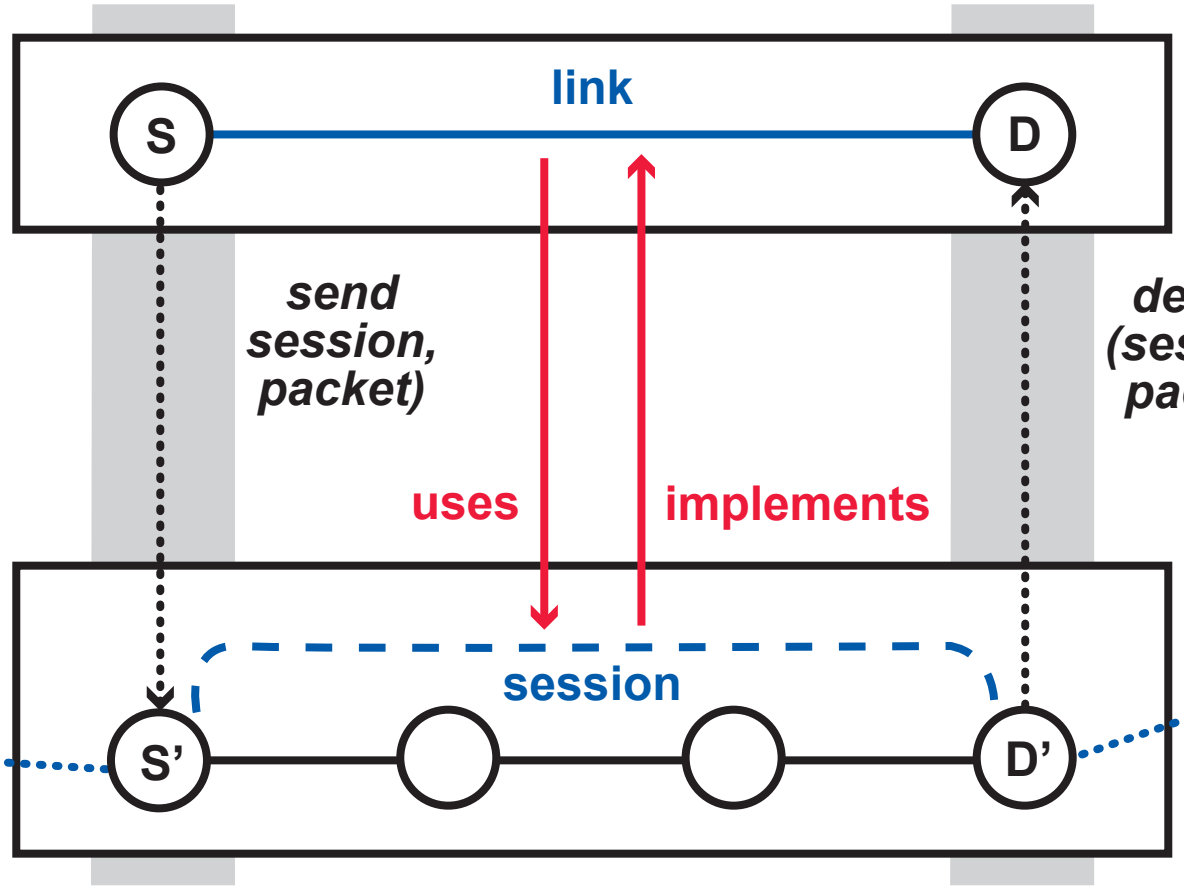

Note that the interface between networks is the same as the user interface to a network.

Every network has the same interface, above and below—this makes networks composable like Lego blocks.

and it only works because sessions are parts of networks

COMPOSITION OPERATORS ON NETWORKS 3

LAYERING in the Internet often looks like this.

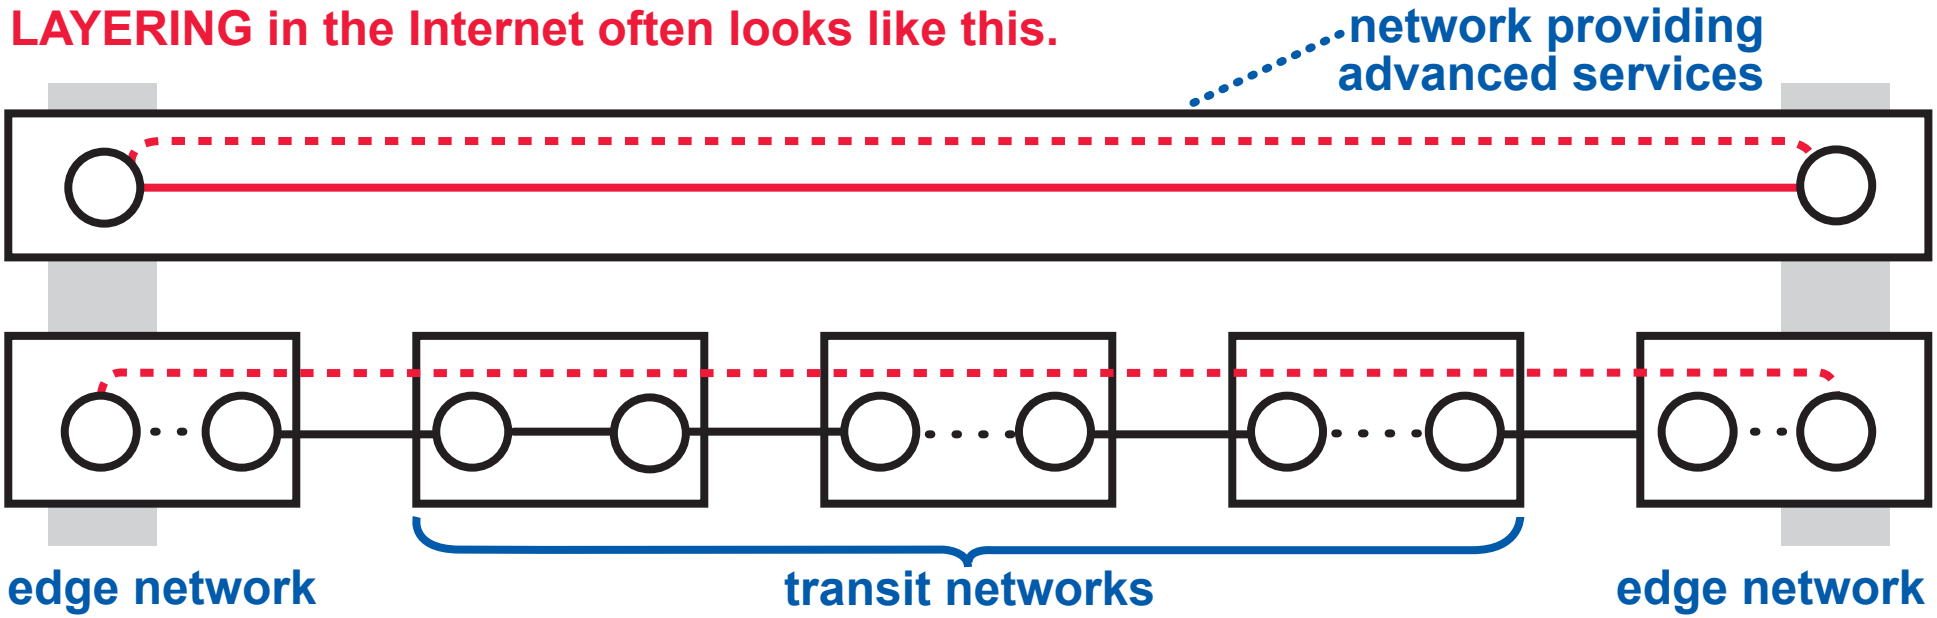

COMPOSITION OPERATORS ON NETWORKS 3

LAYERING in the Internet often looks like this.

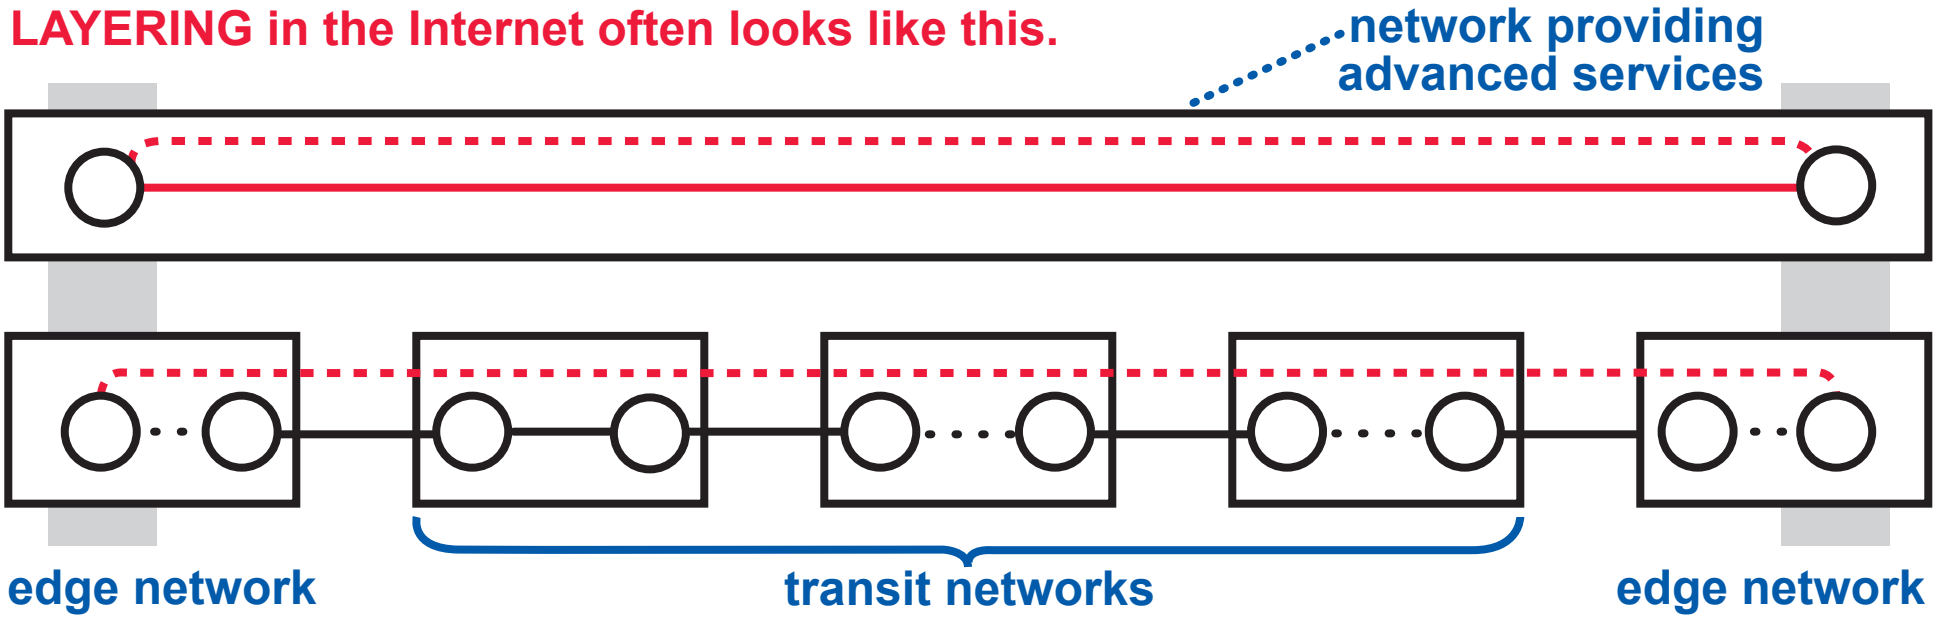

SUBDUCTION: The enabler of innovation and evolution! Allows users of a service to communicate with machines that don't use the service . . . until its popularity spreads.

formally, subduction is a combination of bridging and layering that preserves the properties of both

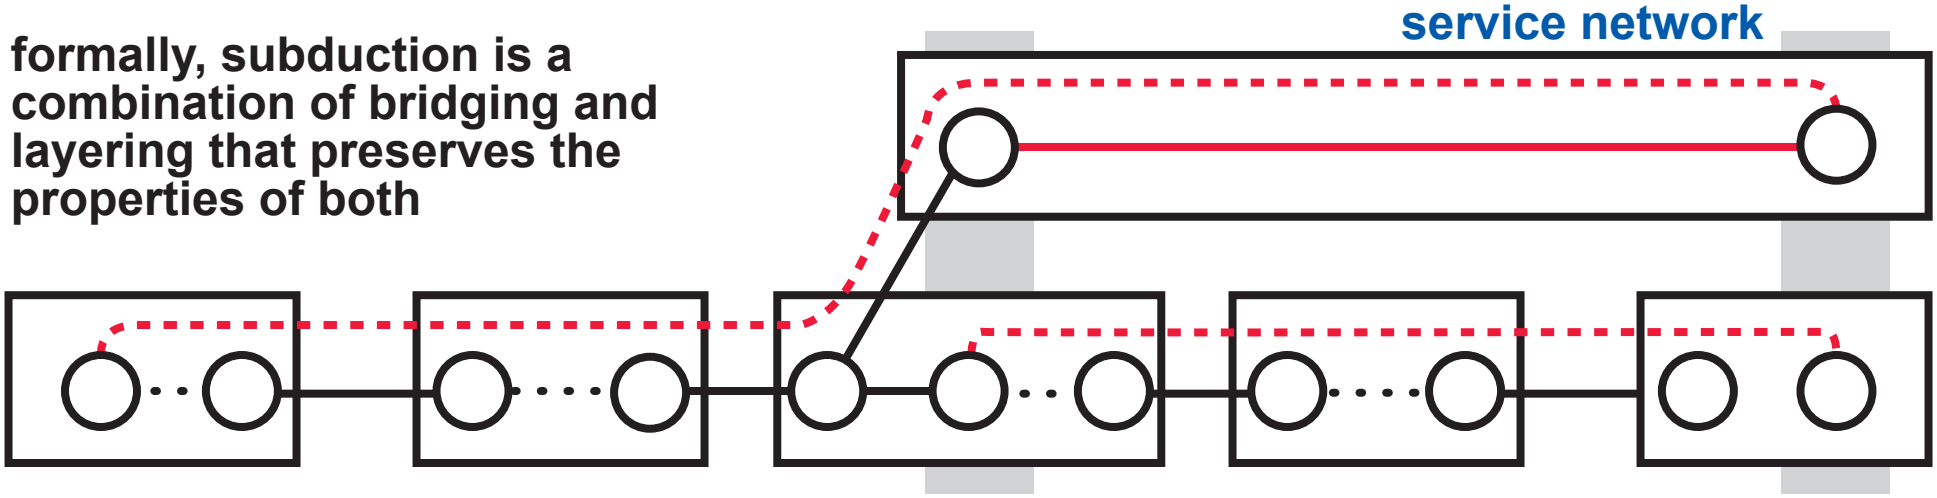

THE REAL INTERNET ARCHITECTURE

THE FLAT INTERNET:

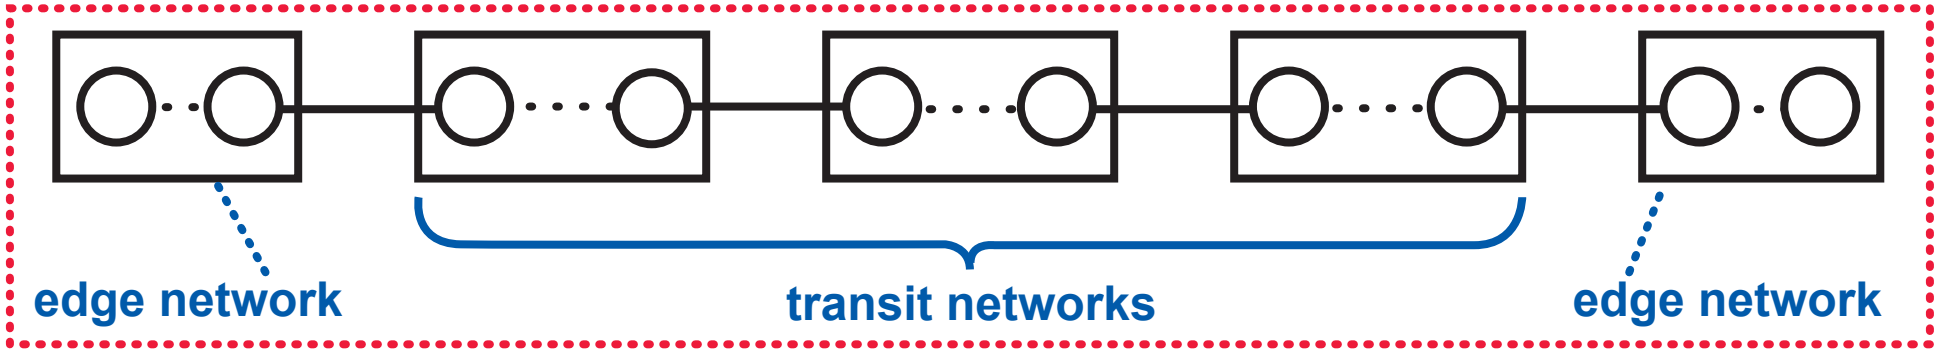

In the real Internet architecture, dozens of network types . . .

. . . with different designs and purposes . . .

. . . are composed (with layering and subduction) both above and below the flat Internet.

PATTERNS OBSERVED	PURPOSE	COMPOSITION
	to span multiple, heterogeneous networks	one overlay, multiple underlays
	to improve the scalability and flexibility of routing	one overlay, one underlay
	to share or "slice" resources	multiple overlays, one underlay
	to provide enhanced network services	end-to-end overlays on the flat Internet

this is the only one in the classic architecture

THE OLD END-TO-END PRINCIPLE:

Functions of the Internet should be minimized, so basic service is efficient, and no one pays for services they don't use.

user machine is outside the network

user interface is link or NIC

middleboxes are bad

there is global routing in the Internet, local routing in the link layer, and no other routing

THE OLD END-TO-END PRINCIPLE:

Functions of the Internet should be minimized, so basic service is efficient, and no one pays for services they don't use.

user machine is outside the network

user interface is link or NIC

middleboxes are bad

there is global routing in the Internet, local routing in the link layer, and no other routing

THE NEW END-TO-END PRINCIPLE

Functions of the FLAT INTERNET should be minimized . . .

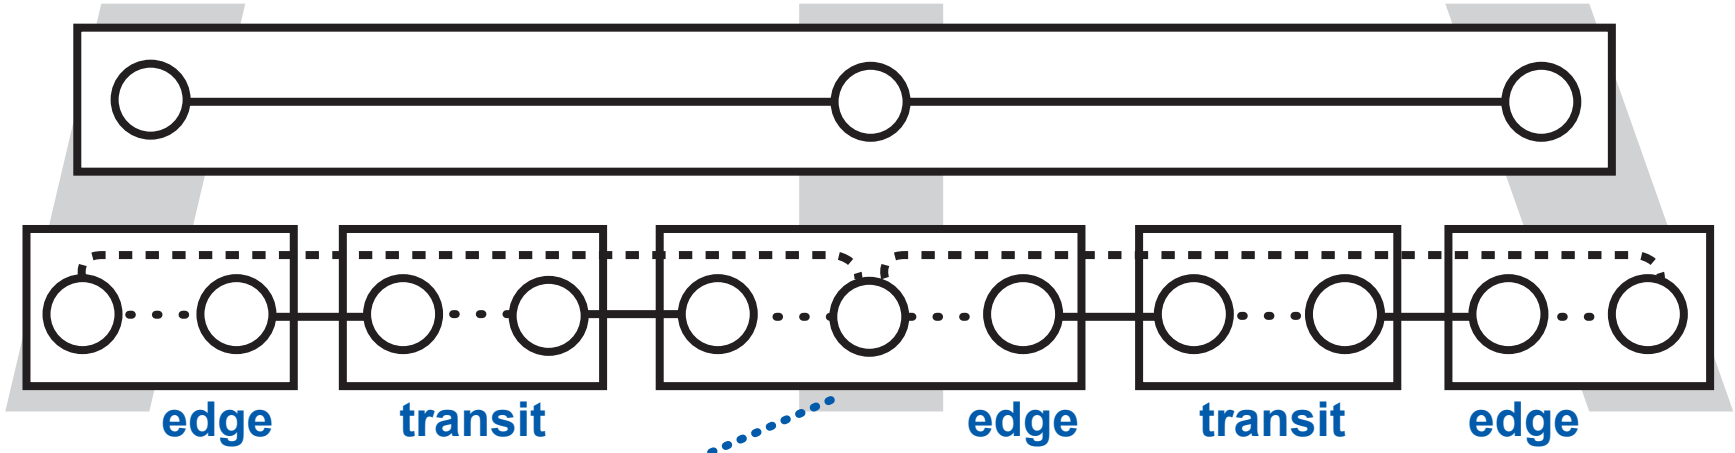

user machines host members of the flat Internet

this machine hosts a user member of the flat Internet and a middlebox (infrastructure member) of the overlay

any module (network) can have routing, and many do

SECURITY IN AN ENTERPRISE NETWORK

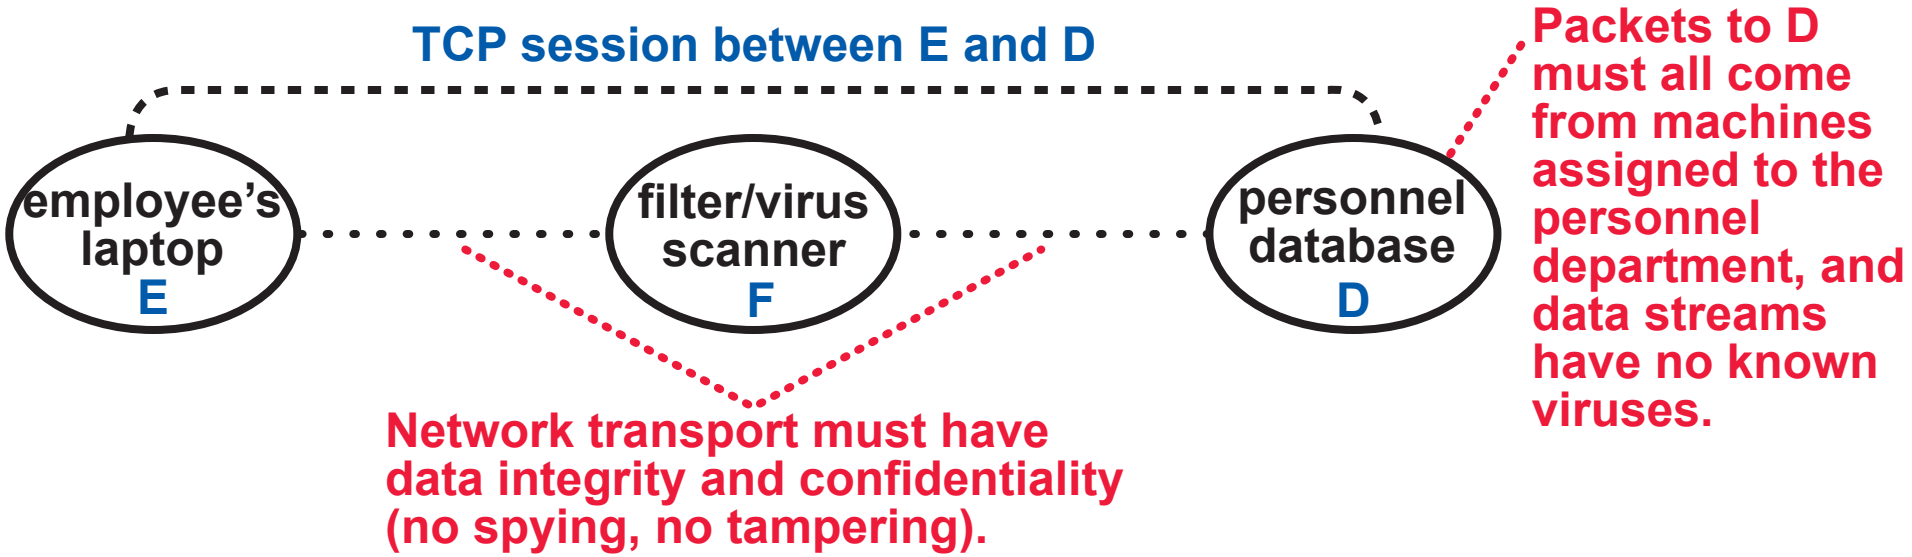

LEMMAS

- The filter/scanner checks that packet source names assigned to personnel, scans byte streams for viruses. — *verify the component*
- All packets of a session to D must pass through the same instance of the filter/scanner. — *verify network routing*
- If a packet has a source name assigned to personnel, then it was sent by a machine assigned to the personnel department. — *?*
- Packet headers are not altered during transport (over a chain of links, forwarders, and middleboxes). — *?*

COMPOSITIONAL NETWORK ARCHITECTURE

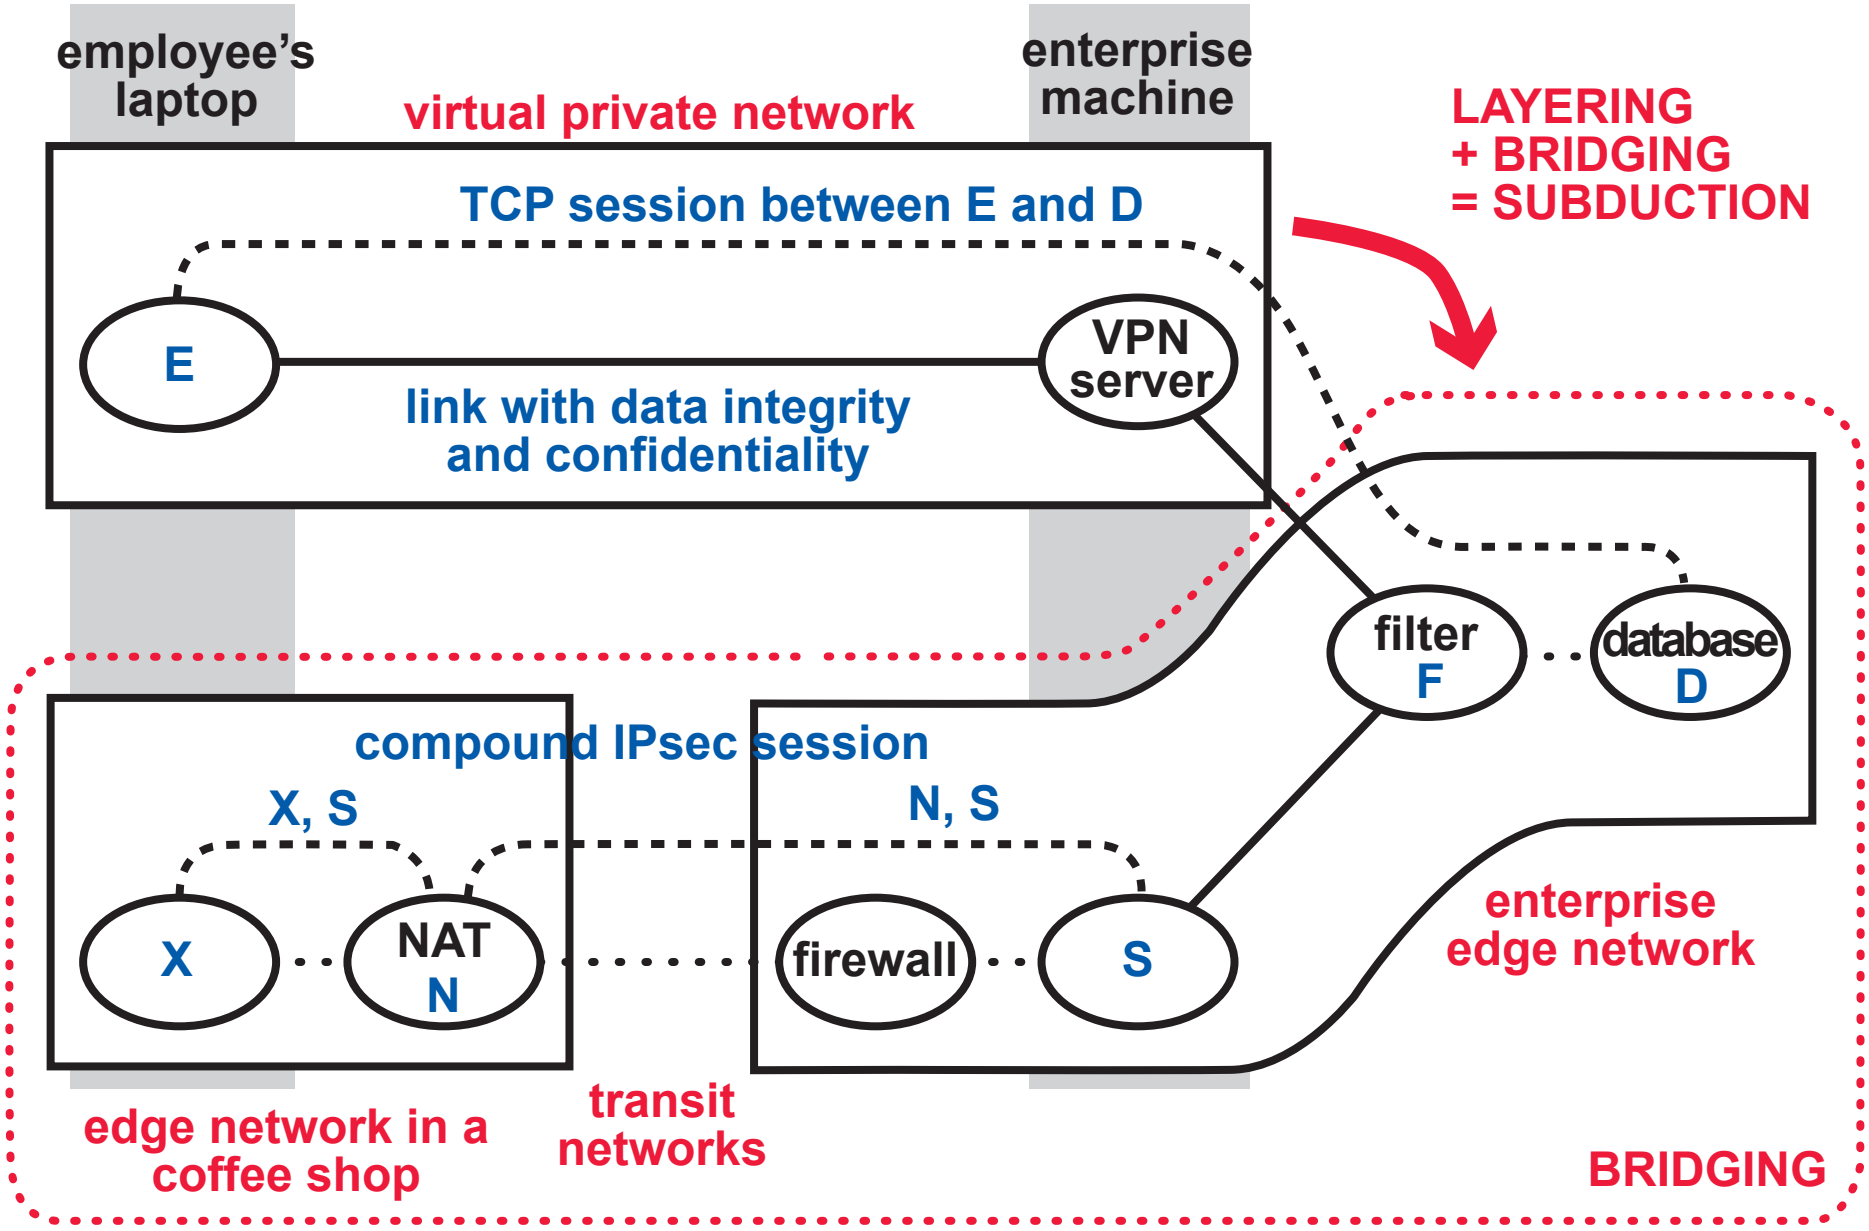

VERIFICATION OF THE LEMMAS

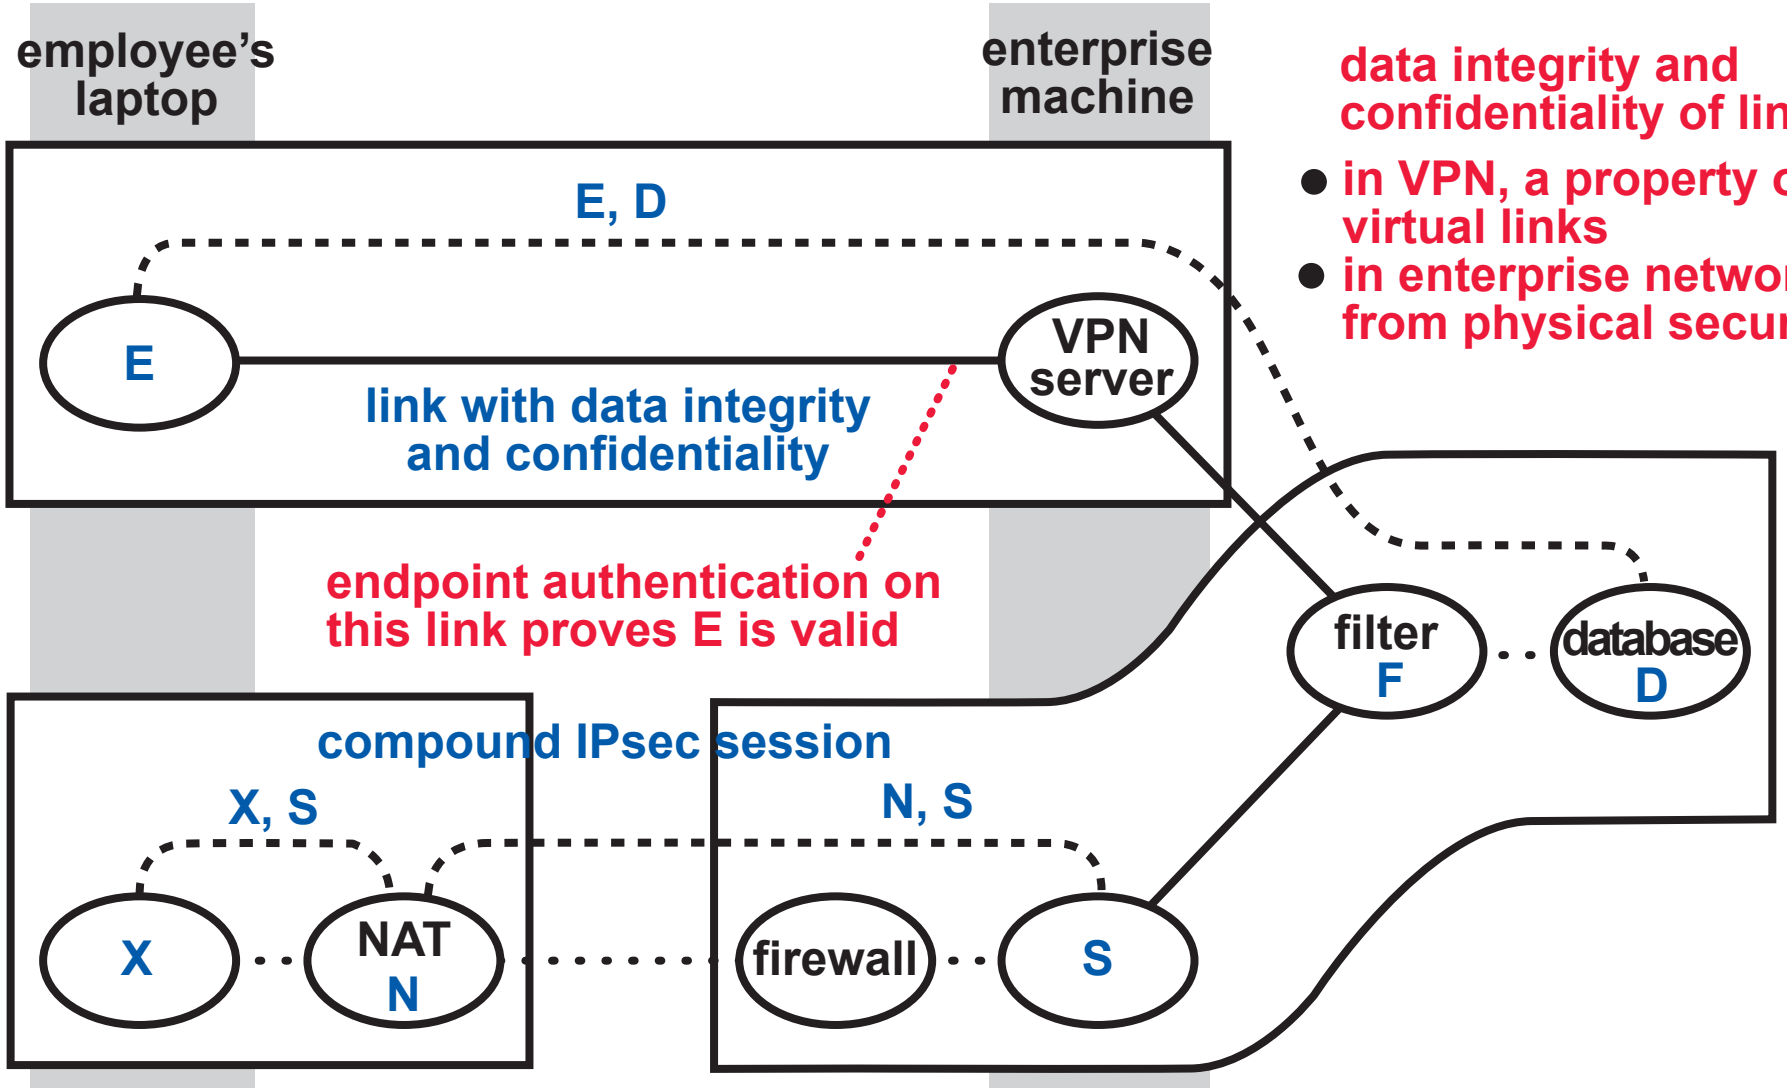

data integrity and confidentiality of links:

- in VPN, a property of virtual links
- in enterprise network, from physical security

packet headers are not altered during transport:

- verification of middleboxes
- packet headers not altered by forwarding

this is an intrinsic property of networks in the model— packet rewriting is replaced by the formal semantics of layering

VERIFICATION SUMMARY

STRUCTURE

- properties are attached to members, links, paths, sessions
- every property applies within one network (or bridged networks)
- some properties are guaranteed by structure alone

the structure can be implemented and verified once, then reused (we have a P4 implementation)

these are the same network, or an overlay/underlay pair

MODULARITY

- every property is proved within one network (or bridged networks)
- this can be a huge improvement in scalability, as the number of forwarding rules in combined networks is often combinatoric

VARIED TECHNIQUES

- proof of cryptographic algorithms
- proof of component programs (hardware and/or software)
- analysis of routing and forwarding (“network verification”)

the full story . . .

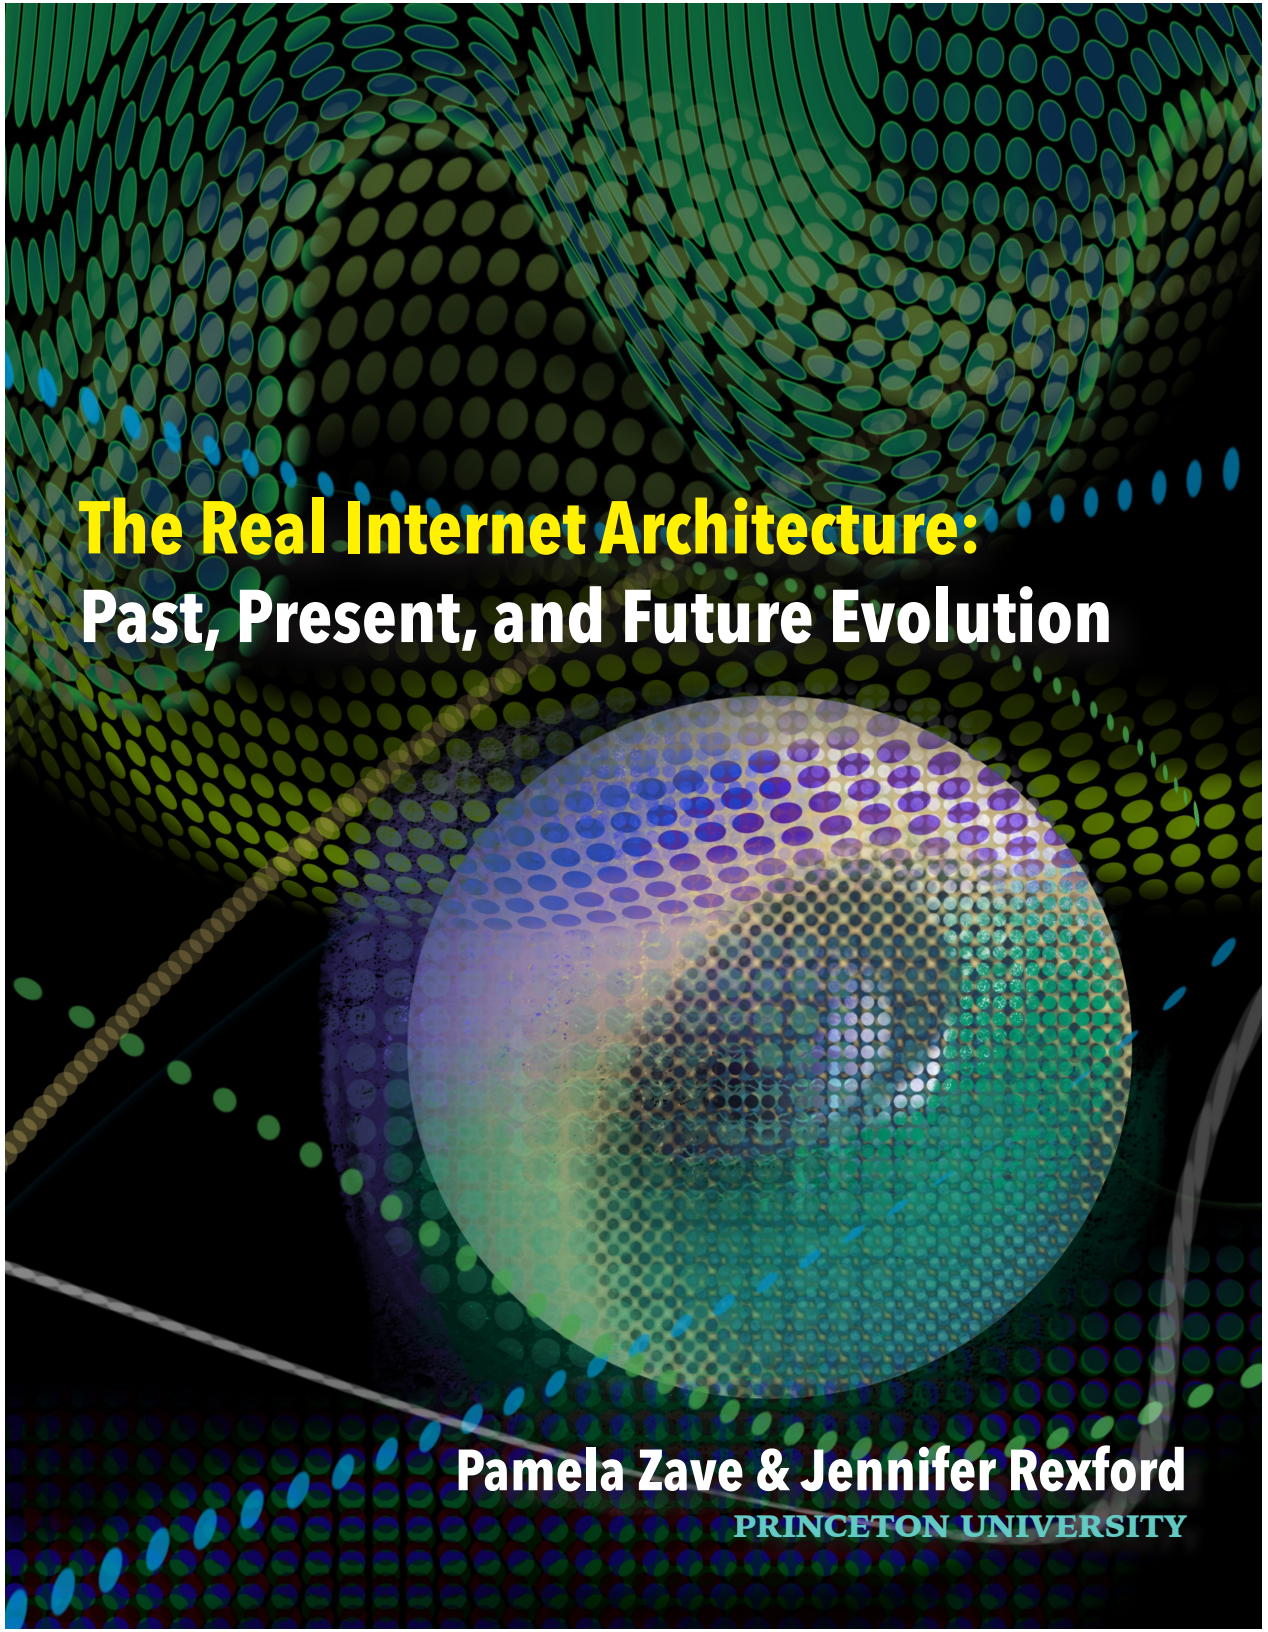

WHAT CAN YOU FIND IN THIS BOOK?

- PATTERNS
- NEW SOLUTIONS TO OLD PROBLEMS
- A FORMAL MODEL

HOW CAN YOUR RESEARCH BENEFIT FROM THEM?

PATTERNS

A PATTERN IS A RECURRING PROBLEM . . .

. . . AND A FAMILY OF RELATED SOLUTIONS

THE BOOK IS FULL OF PATTERNS . . .

. . . because the same problems recur in every network, and the new model exposes the repetition

- **new namespaces**
- **flexible and scalable routing**
- **resource slicing**
- **security and privacy**
- **performance optimization**
- **mobility**
- **multicast**

PATTERNS are a better way to teach networking

- **meaningful and memorable**
- **emphasize engineering decisions: which solution for *this* instance of the problem?**
- **concise when you need them to be concise**

PATTERNS can expand the applicability of your research

- **where else in a network architecture do the same problems arise?**

NEW SOLUTIONS TO OLD PROBLEMS

The compositional model expands

design freedom.

YOU may be the one to solve stubborn

practical problems by seeing them

in a new way.

A FORMAL MODEL OF

COMPOSITIONAL NETWORK ARCHITECTURE

AN OPPORTUNITY FOR COMMUNITY EFFORT AND PROGRESS

- ensure compatibility and inter-operation of results
- develop coordinated tools

  possibly based on P4
- build on each other's work

FOR MORE INFORMATION

DRAFTS OF THE BOOK WILL BE AVAILABLE VERY SOON

. . . and sent to interested readers for comments

RELATED WORK

- Pamela Zave and Jennifer Rexford
  "The compositional architecture of the Internet"
  Communications of the ACM
  March 2019

- "Patterns and interactions in network security"
  ACM Computing Surveys
  December 2020

- “The design space of network mobility"
  Recent Advances in Networking
  ACM SIGCOMM ebook
  edited by Olivier Bonaventure and Hamed Haddadi
  2013